

Single-Photon Depth Imaging Using a Union-of-Subspaces Model

Donggeek Shin, Jeffrey H. Shapiro, and Vivek K Goyal

Abstract—Light detection and ranging systems reconstruct scene depth from time-of-flight measurements. For low light-level depth imaging applications, such as remote sensing and robot vision, these systems use single-photon detectors that resolve individual photon arrivals. Even so, they must detect a large number of photons to mitigate Poisson shot noise and reject anomalous photon detections from background light. We introduce a novel framework for accurate depth imaging using a small number of detected photons in the presence of an unknown amount of background light that may vary spatially. It employs a Poisson observation model for the photon detections plus a union-of-subspaces constraint on the discrete-time flux from the scene at any single pixel. Together, they enable a greedy signal-pursuit algorithm to rapidly and simultaneously converge on accurate estimates of scene depth and background flux, without any assumptions on spatial correlations of the depth or background flux. Using experimental single-photon data, we demonstrate that our proposed framework recovers depth features with 1.7 cm absolute error, using 15 photons per image pixel and an illumination pulse with 6.7-cm scaled root-mean-square length. We also show that our framework outperforms the conventional pixelwise log-matched filtering, which is a computationally-efficient approximation to the maximum-likelihood solution, by a factor of 6.1 in absolute depth error.

Index Terms—Computational imaging, LIDAR, single-photon imaging, union-of-subspaces, greedy algorithms.

I. INTRODUCTION

A conventional light detection and ranging (LIDAR) system, which uses a pulsed light source and a single-photon detector, forms a depth image pixelwise using the histograms of photon detection times. The acquisition times for such systems are made long enough to detect hundreds of photons per pixel for the finely binned histograms these systems require to do accurate depth estimation. In this letter, we introduce a framework for accurate depth imaging using only a small number of photon detections per pixel, despite the presence of an unknown amount of spatially-varying background light in the scene. Our framework uses a Poisson observation model for the photon detections plus a union-of-subspaces constraint on the scene’s discrete-time flux at any single pixel. Using a greedy signal-pursuit algorithm—a modification of CoSaMP [1]—we solve for accurate estimates of scene depth and background flux. Our method forms estimates pixelwise

and thus avoids assumptions on transverse spatial correlations that may hinder the ability to resolve very small features. Using experimental single-photon data, we demonstrate that our proposed depth imaging framework outperforms log-matched filtering, which is the maximum-likelihood (ML) depth estimator given zero background light.

Because our proposed framework is photon-efficient while using an estimator that is pixelwise and without background calibration, it can be useful for dynamic low light-level imaging applications, such as environmental mapping using unmanned aerial vehicles.

A. Prior Art

The conventional LIDAR technique of estimating depth using histograms of photon detections is accurate when the number of photon detections is high. In the low photon-count regime, the depth solution is noisy due to shot noise. It has been shown that image denoising methods, such as wavelet thresholding, can boost the performance of scene depth recovery in the presence of background noise [2]. Also, using an imaging model that incorporates occlusion constraints was proposed to recover an accurate depth map [3]. However, these denoising algorithms implicitly assume that the observations are Gaussian distributed. Thus, at low photon-counts, where depth estimates are highly non-Gaussian [4], their performance degrades significantly [5].

First-photon imaging (FPI) [6] is a framework that allows high-accuracy imaging using only the first detected photon at every pixel. It demonstrated that centimeter-accurate depth recovery is possible by combining the non-Gaussian statistics of first-photon detection with spatial correlations of natural scenes. The FPI framework uses an imaging setup that includes a raster-scanning light source and a lensless single-photon detector. More recently, photon-efficient imaging frameworks that use a detector array setup, in which every pixel has the same acquisition time, have also been proposed [5], [7], [8].

We observe two common limitations that exist in the prior active imaging frameworks for depth reconstruction.

- **Over-smoothing:** Many of the frameworks assume spatial smoothness of the scene to mitigate the effect of shot noise. In some imaging applications, however, it is important to capture fine spatial features that only occupy a few image pixels. Using methods that assume spatial correlations may lead to erroneously over-smoothed images that wash out the scene’s fine-scale features. In such scenarios, a robust *pixelwise* imager is preferable.

This material is based upon work supported in part by a Samsung Scholarship, by the US National Science Foundation under Grant No. 1422034, and by the MIT Lincoln Laboratory Advanced Concepts Committee.

D. Shin and J. H. Shapiro are with the Department of Electrical Engineering and Computer Science and the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

V. K. Goyal is with the Department of Electrical and Computer Engineering, Boston University, Boston, MA 02215 USA.

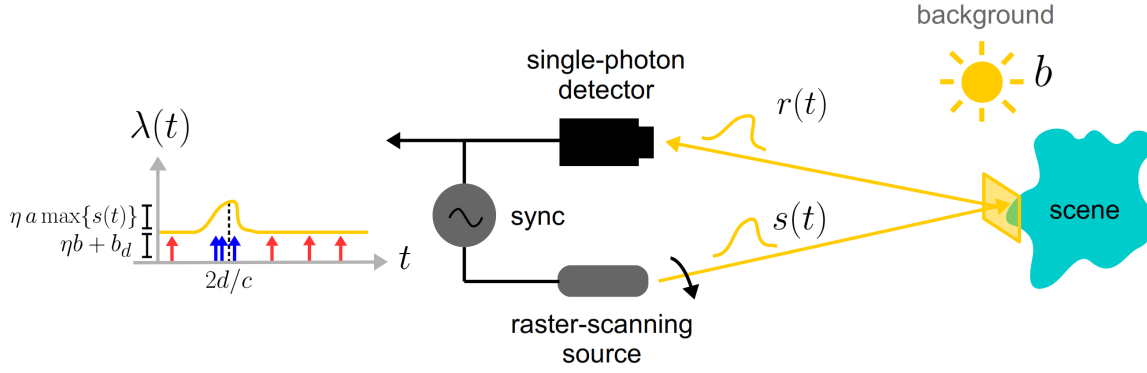


Fig. 1. An illustration of the single-photon imaging setup for one illumination pulse. A pulsed optical source illuminates a scene pixel with photon-flux waveform $s(t)$. The flux waveform $r(t)$ that is incident on the detector consists of the pixel return $as(t - 2d/c)$ —where a is the pixel reflectivity, d is the pixel depth, and c is light speed—plus the background-light flux b . The rate function $\lambda(t)$ driving the photodetection process equals the sum of the pixel return and background flux, scaled by the detector efficiency η , plus the detector's dark-count rate b_d . The record of detection times from the pixel return (or background light plus dark counts) is shown as blue (or red) spikes, generated by the Poisson process driven by $\lambda(t)$.

- **Calibration:** Many imaging methods assume a calibration step to measure the amount of background flux existing in the environment. This calibration mitigates bias in the depth estimate caused by background-photon or dark-count detections, which have high temporal variance. In practical imaging scenarios, however, the background response varies in time, and continuous calibration may not be practical. Furthermore, many methods assume background flux does not vary spatially. Thus, a *calibrationless* imager that performs simultaneous estimation of scene parameters and spatially-varying background flux from photon detections is useful.

In this letter, we propose a novel framework for depth acquisition that is applied pixelwise and without calibration. At each pixel, our imager estimates the background response along with scene depth from photon detections. Similar to [3], we use a union-of-subspaces constraint for modeling the scene parameters. However, our union-of-subspace constraint is defined for both the incoherent signal and background waveform parameters that generate photon detections; the constraint in [3] is defined for only the coherent signal waveform that is perturbed by Gaussian noise, not photon noise.

Using the derived imaging model, we propose a greedy signal pursuit algorithm that accurately solves for the scene parameters at each pixel. We evaluate the photon efficiency of this framework using experimental single-photon data. In the presence of strong background light, we show that our pixelwise imager gives an absolute depth error that is 6.1 times lower than that of the pixelwise log-matched filter.

II. SINGLE-PHOTON IMAGING SETUP

Figure. 1 illustrates our imaging setup, for one illumination pulse, in which photon detections are made. A focused optical source, such as a laser, illuminates a pixel of the scene with the pulse waveform $s(t)$ that starts at time 0 and has root-mean-square pulsewidth T_p . This illumination is repeated every T_r seconds for a sequence of N_s pulses. The single-photon detector, in conjunction with a time correlator, is used to time stamp individual photon detections, relative to the time at which the

immediately preceding pulse was transmitted. These detection times, which are observations of a time-inhomogeneous Poisson process, whose rate function combines contributions from pixel return, background light, and dark counts, are used to estimate scene depth for the illuminated pixel. This pixelwise acquisition process is repeated for $N_x \times N_y$ image pixels by raster scanning the light source in the transverse directions.

III. FORWARD IMAGING MODEL

In this section, we study the relationship between the photon detections and the scene parameters. For simplicity of exposition and notation, we focus on one pixel; this is repeated for each pixel of a raster-scanning or array-detection setup.

Let a , d , and b be unknown scalar values that represent reflectivity, depth, and background flux at the given pixel. The reflectivity value includes the effects of radial fall-off, view angle, and material properties. Then, after illuminating the scene pixel with a single pulse $s(t)$, the backreflected waveform that is incident at the single-photon detector is

$$r(t) = as(t - 2d/c) + b, \quad t \in [0, T_r]. \quad (1)$$

A. Photodetection statistics

Using (1), we observe that the rate function that generates the photon detections is

$$\lambda(t) = \eta(as(t - 2d/c) + b) + b_d, \quad t \in [0, T_r], \quad (2)$$

where $\eta \in (0, 1]$ is the quantum efficiency of the detector and $b_d \geq 0$ is the dark-count rate of the single-photon detector.

Let Δ be the time bin duration of the single-photon detector. Then, we define $M = T_r/\Delta$ to be the total number of time bins that capture photon detections. Let \mathbf{y} be the vector of size $M \times 1$ that contains the number of photon detections at each time bin after we illuminate the pixel N_s times with pulse waveform $s(t)$. Then, from photodetection theory [9], we can derive

$$\mathbf{y}_k \sim \text{Poisson} \left(N_s \int_{(k-1)\Delta}^{k\Delta} [\eta(as(t - 2d/c) + b) + b_d] dt \right), \quad (3)$$

for $k = 1, \dots, M$. Note that we have assumed that our total pixelwise acquisition time $N_s T_r$ is short enough that b is constant during that time. Let

$$\mathbf{v}_j = \int_{(j-1)\epsilon}^{j\epsilon} a\delta(t - 2d/c) dt, \quad (4)$$

$$\mathbf{S}_{i,j} = \int_{(i-1)\Delta}^{i\Delta} \int_{(j-1)\epsilon}^{j\epsilon} N_s \eta s(t - y) dt dy, \quad (5)$$

$$B = N_s \Delta (\eta b + b_d), \quad (6)$$

for $i = 1, \dots, M$, $j = 1, \dots, N$, where ϵ is a small number, such that T_r is divisible by ϵ and $N = T_r/\epsilon$. Defining $\mathbf{1}_{M \times 1}$ to be an $M \times 1$ vector of ones, we can approximate the rate function in (3) and rewrite the distribution as

$$\mathbf{y}_k \sim \text{Poisson}((\mathbf{S}\mathbf{v} + B\mathbf{1}_{M \times 1})_k), \quad (7)$$

for $k = 1, \dots, M$. Finally, defining $\mathbf{A} = [\mathbf{S}, \mathbf{1}_{M \times 1}]$ and $\mathbf{x} = [\mathbf{v}^T, B]^T$, we can further rewrite (7) as

$$\mathbf{y}_k \sim \text{Poisson}((\mathbf{A}\mathbf{x})_k). \quad (8)$$

So far, we have simplified the pixelwise single-photon observation model, such that the photon-count vector $\mathbf{y} \in \mathbb{N}^{M \times 1}$ is a linear measurement of scene response vector $\mathbf{x} \in \mathbb{R}_+^{(N+1) \times 1}$ corrupted by Poisson noise.

B. Scene parameter constraints

Using the expression in (4), we observe that

$$\mathbf{v}_j = \int_{(j-1)\epsilon}^{j\epsilon} a\delta(t - 2d/c) dt \quad (9)$$

$$= a \mathbb{1}_{\{x: (j-1)\epsilon \leq 2x/c < j\epsilon\}}(d), \quad (10)$$

for $j = 1, \dots, N$, where $\mathbb{1}_A(x)$ is an indicator function that equals 1 if $x \in A$ and 0 otherwise. In other words, vector \mathbf{v} has exactly one nonzero element, and the value and index of the nonzero element represents the scene reflectivity and depth at an image pixel, respectively.

We defined our $(N+1) \times 1$ signal \mathbf{x} to be a concatenation of \mathbf{v} , which is the scene response vector of size N , and B , which is the scalar representing background flux. Since \mathbf{v} has exactly one nonzero entry, \mathbf{x} lies in the union of N subspaces defined as

$$\mathcal{S}_N = \bigcup_{k=1}^N \{\mathbf{x} \in \mathbb{R}^{N+1} : \mathbf{x}_{\{1,2,\dots,N\} \setminus \{k\}} = \mathbf{0}\}, \quad (11)$$

where each subspace is of dimension 2.

IV. SOLVING THE INVERSE PROBLEM

Using accurate photodetection statistics and scene constraints, we have interpreted the problem of robust single-photon depth imaging as a noisy linear inverse problem, where the signal of interest \mathbf{x} lies in the union of subspaces \mathcal{S}_N .

Using (8), the observed photon count histogram \mathbf{y} has the probability mass function

$$f_Y(\mathbf{y}; \mathbf{A}, \mathbf{x}) = \prod_{k=1}^M \frac{e^{-(\mathbf{A}\mathbf{x})_k} (\mathbf{A}\mathbf{x})_k^{\mathbf{y}_k}}{\mathbf{y}_k!}. \quad (12)$$

Thus, neglecting terms in the negative log-likelihood function that are dependent on \mathbf{y} but not on \mathbf{x} , we can define the objective function

$$\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y}) = \sum_{k=1}^M [(\mathbf{A}\mathbf{x})_k - \mathbf{y}_k \log(\mathbf{A}\mathbf{x})_k]. \quad (13)$$

This objective can be proved to be convex in \mathbf{x} .

We solve for \mathbf{x} by minimizing $\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y})$ with the constraint that \mathbf{x} lies in the union of subspaces \mathcal{S}_N . Also, because photon flux is a non-negative quantity, the minimization results in a more accurate estimate when we include a non-negative signal constraint. In summary, the optimization problem that we want to solve can be written as

$$\begin{aligned} & \underset{\mathbf{x}}{\text{minimize}} && \mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y}) \\ & \text{s.t.} && \mathbf{x} \in \mathcal{S}_N, \\ & && \mathbf{x}_i \geq 0, \quad i = 1, \dots, (N+1). \end{aligned} \quad (14)$$

To solve our constrained optimization problem, we propose an algorithm that is inspired by existing fast greedy algorithms for sparse signal pursuit. CoSaMP [1] is a greedy algorithm that finds a K -sparse approximate solution to a linear inverse problem. We modify the CoSaMP algorithm so that we obtain for a solution constrained to the union of subspaces \mathcal{S}_N , instead of a K -sparse one.

Algorithm 1 Single-photon depth imaging using a union-of-subspaces model

Input: $\mathbf{y}, \mathbf{A}, \delta$
Output: $\mathbf{x}^{(k)}$
Initialize $\mathbf{x}^{(0)} \leftarrow \vec{0}$, $\mathbf{u} \leftarrow \mathbf{y}$, $k \leftarrow 0$;
repeat
 $k \leftarrow k + 1$;
 $\hat{\mathbf{x}} \leftarrow \mathbf{A}^T \mathbf{u}$;
 $\Omega \leftarrow \text{supp}((\hat{\mathbf{x}}_{1:N})_{[1]}) \cup \text{supp}(\mathbf{x}_{1:N}^{(k-1)}) \cup \{N+1\}$;
 $\mathbf{b}|_{\Omega} \leftarrow \mathbf{A}_{\Omega}^{\dagger} \mathbf{y}$; $\mathbf{b}|_{\Omega^c} \leftarrow \mathbf{0}$;
 $\mathbf{x}^{(k)} \leftarrow \mathcal{T}_0 \left([(\mathbf{b}_{1:N})_{[1]}^T, \mathbf{b}_{N+1}]^T \right)$
 $\mathbf{u} \leftarrow \mathbf{y} - \mathbf{A}\mathbf{x}^{(k)}$
until $\|\mathbf{x}^{(k-1)} - \mathbf{x}^{(k)}\|_2^2 < \delta$

Our proposed greedy algorithm is given in Algorithm 1. We define $\mathcal{T}_0(\mathbf{x})$ to be the thresholding operator setting all negative entries of \mathbf{x} to zero, $\text{supp}(\mathbf{x})$ to be the set of indices of nonzero elements of \mathbf{x} , and $\mathbf{x}_{[k]}$ to be the vector that approximates \mathbf{x} with its k largest terms. Also, we take \mathbf{A}_S to be a matrix with columns of \mathbf{A} chosen by the index set S . Finally, we use \mathbf{A}^T and \mathbf{A}^{\dagger} to denote the transpose and pseudo-inverse of matrix \mathbf{A} , respectively.

To solve the constrained optimization problem in (14), our algorithm iteratively performs the following steps:

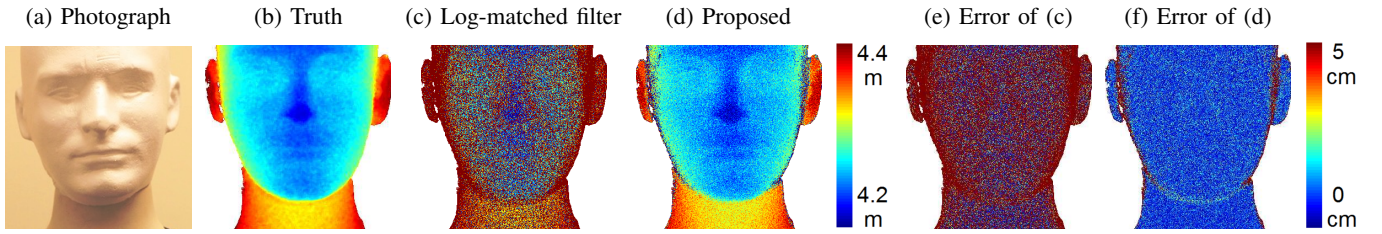


Fig. 2. Experimental pixelwise depth imaging results using single photon observations. The number of photon detections at every pixel was set to be 15. The figure shows the (a) photograph of imaged face, (b) ground-truth depth, (c) depth from log-matched filtering, which is approximately ML and (d) depth using our method. Also, (e) and (f) show the absolute depth error maps for ML and our framework, respectively.

- 1) gradient descent on $\mathcal{L}(\mathbf{x}; \mathbf{A}, \mathbf{y})$, which is approximated by the squared ℓ_2 -norm $\|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2$ for computational efficiency;
- 2) projection of the intermediate estimate onto the closest subspace in the union of subspaces \mathcal{S}_N ; and
- 3) projection of the intermediate estimate onto the non-negative cone,

until a convergence criterion is satisfied. We define convergence of the solution as $\|\mathbf{x}^{(k-1)} - \mathbf{x}^{(k)}\|_2^2 < \delta$, where δ is a small number.

V. EXPERIMENTAL RESULTS

To validate our imaging framework, we used a dataset collected by D. Venkatraman for the First-Photon Imaging project [6]; this dataset and others are available from [10]. The experimental setup uses a pulsed laser diode with pulsewidth $T_p = 270$ ps and repetition period $T_r = 100$ ns. A two-axis galvo was used to scan 350×350 pixels of a mannequin face at a distance of about 4 m. A lensless SPAD detector with quantum efficiency $\eta = 0.35$ was used for detection. The background light level was set using an incandescent lamp. The original mannequin data from [10] had the background count rate approximately equal to the signal count rate. Our experiment uses the cropped data showing only the mannequin's face, where the background count rate was approximately 0.1 of the average signal count rate. Although we used a raster-scanning setup for our single-photon imaging experiments, since our imaging algorithm is applied pixelwise, it can be also used for imaging with a floodlight illumination source and a single-photon detector array.

We could compare our imaging method with the ML estimator of scene parameters. Unfortunately, due to nonzero background flux, ML estimation of a , b , and d requires minimizing a non-convex cost function, leading to a solution without an accuracy guarantee. Thus, zero background is assumed conventionally such that the ML depth estimate reduces to the simple log-matched filter [11]:

$$\hat{d}_{\text{ML}} = \frac{\epsilon}{2} \left(\arg \max_{i \in \{1, \dots, n\}} \log \mathbf{S}_i^T \mathbf{y} \right). \quad (15)$$

We use (15) as the baseline depth estimator that is compared with our proposed estimator.

Figure 2 shows the results of recovering depth of the mannequin face using single-photon observations. The kernel matrix \mathbf{S} was obtained by an offline measurement of the pulse shape. Note that this measurement depends only on the source,

not on properties of the scene. The ground-truth depth, shown in Fig. 2(b), was generated separately by using background-calibrated ML estimation from 200 photons at each pixel.

In our depth imaging experiment, the number of photon detections at each pixel was set to 15. We observe that, due to extraneous background photon detections, the log-matched filter estimate in Fig. 2(c) (average absolute error = 10.3 cm) is corrupted with high-variance noise and the facial features of the mannequin are not visible. On the other hand, our estimate, shown in Fig. 2(d), shows high-accuracy depth recovery (average absolute error = 1.7 cm). As shown by the error maps in Fig. 2(e), (f), both methods fail in depth recovery in the face boundary regions, where very little light is reflected back from the scene to the single-photon detector and the signal-to-background ratio is thus very low. Also, we observe that our estimated average background level over all pixels was $\hat{B} = 1.4 \times 10^{-3}$, which is very close to the calibrated true background level $B = 1.3 \times 10^{-3}$. In this experiment, we had $M = N = 801$. Also, we set $\delta = 10^{-4}$ and the average number of iterations until convergence was measured to be 2.1 over all pixels. Code and data used to generate results can be downloaded from [12].

VI. CONCLUSIONS AND FUTURE WORK

In this letter, we presented an imaging framework for calibrationless, pixelwise depth reconstruction using single-photon observations. Our imaging model combined photon detection statistics with the discrete-time flux constraints expressed using a union-of-subspaces model. Then, using our imaging model, we developed a greedy algorithm that recovers scene depth by solving a constrained optimization problem.

Our pixelwise imaging framework can be used in low light-level imaging applications, where the scene being imaged has fine features and filtering techniques that exploit patchwise smoothness can potentially wash out those details. For example, it can be useful in applications such as airborne remote sensing [13], where the aim is to recover finely-featured 3D terrain maps.

It is straightforward to generalize the proposed single-photon imaging framework to multiple-depth estimation, where more than one reflector may be present at each pixel. In the case of estimating depths of K reflectors at a pixel, the 1-sparsity assumption must be changed to more general K -sparsity assumption when defining the union-of-subspaces constraint.

REFERENCES

- [1] D. Needell and J. A. Tropp, “CoSaMP: Iterative signal recovery from incomplete and inaccurate samples,” *Appl. Comput. Harmon. Anal.*, vol. 26, no. 3, pp. 301–321, 2009.
- [2] B.-Y. Sun, D.-S. Huang, and H.-T. Fang, “Lidar signal denoising using least-squares support vector machine,” *IEEE Signal Process. Lett.*, vol. 12, no. 2, pp. 101–104, 2005.
- [3] P. T. Boufounos, “Depth sensing using active coherent illumination,” in *Proc. IEEE Int. Conf. Acoust., Speech, and Signal Process.*, 2012, pp. 5417–5420.
- [4] A. McCarthy, R. J. Collins, N. J. Krichel, V. Fernández, A. M. Wallace, and G. S. Buller, “Long-range time-of-flight scanning sensor based on high-speed time-correlated single-photon counting,” *Appl. Optics*, vol. 48, no. 32, pp. 6241–6251, 2009.
- [5] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, “Photon-efficient computational 3-d and reflectivity imaging with single-photon detectors,” *IEEE Trans. Comput. Imaging*, vol. 1, 2015, to appear.
- [6] A. Kirmani, D. Venkatraman, D. Shin, A. Colaço, F. N. Wong, J. H. Shapiro, and V. K. Goyal, “First-photon imaging,” *Science*, vol. 343, no. 6166, pp. 58–61, 2014.
- [7] D. Shin, A. Kirmani, V. K. Goyal, and J. H. Shapiro, “Computational 3d and reflectivity imaging with high photon efficiency,” in *Proc. IEEE Int. Conf. Image Process.*, Oct. 2014, pp. 46–50.
- [8] Y. Altmann, X. Ren, A. McCarthy, G. S. Buller, and S. McLaughlin, “Lidar waveform based analysis of depth images constructed using sparse single-photon data,” arXiv preprint arXiv:1507.02511, 2015.
- [9] D. L. Snyder, *Random Point Processes*. Wiley, New York, 1975.
- [10] “GitHub repository for photon-efficient imaging,” <https://github.com/photon-efficient-imaging/sample-data/>.
- [11] R. E. Blahut, *Principles and Practice of Information Theory*. Addison-Wesley Longman Publishing Co., Inc., 1987.
- [12] “GitHub repository for union-of-subspace imaging,” <https://github.com/photon-efficient-imaging/uos-imaging/>.
- [13] M. Nilsson, “Estimation of tree heights and stand volume using an airborne lidar system,” *Remote Sensing of Environment*, vol. 56, no. 1, pp. 1–7, 1996.